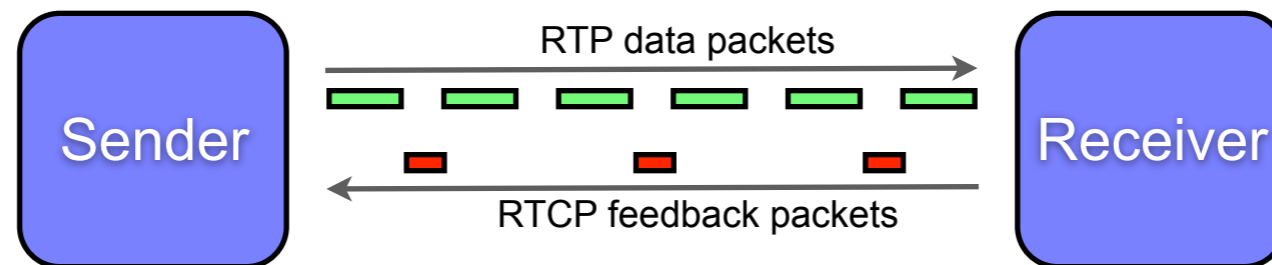


Use of ECN with RTP Circuit Breaker

draft-ietf-avtcore-rtp-circuit-breakers-16

Colin Perkins

Review: RTP Congestion Control



- Dynamics of congestion control depend on rate of feedback, speed of adaptation, and feedback data
- RTP does each of these very differently to TCP
 - Feedback and adaptation done on timescale matching video frame rate, not RTT – for low RTT flows, may report/adapt once per several RTT
 - Cannot react to ECN-CE or loss on timescale of RTT; reacts to delayed aggregated feedback → nominal rate = $f(\text{loss rate}, \text{marking rate})$
 - Sending rate is bursty, but not greedy – there is a maximum usable rate
 - Low latency critical; packet loss very disruptive to user experience

RTP Circuit Breaker

- **Fast-trip circuit breaker**
 - Media quality – data received, but quality too poor to be useful
 - Media timeout – feedback shows data not being received
 - RTCP timeout – feedback not being returned
 - Congestion – RTCP congestion feedback shows persistent excessive congestion on the path
 - Operates on same feedback used for congestion control
 - Persistent → minimum trigger interval = $\max(3 \times \text{RTCP interval}, 10 \times \text{RTT}, 10 \times \text{GoP duration})$
 - Excessive → rate $> 10 \times$ TCP equivalent rate derived via Padhye equation ($1/\sqrt{p}$), responded to ECN-CE marks the same as it responded to loss [RFC3168]
- **Conforms to draft-ietf-tsvwg-circuit-breaker-15**

RTP Circuit Breaker and ECN

- Circuit breaker response to ECN-CE controversial
 - Removed ECN from RTP circuit breaker draft, since urgent to publish and ECN not used in practice, but need to resolve for future deployments
 - The draft said “ECN-CE marked packets SHOULD be treated as if they were lost for the purposes of congestion control”
 - RFC 3168: “Upon the receipt by an ECN-Capable transport of a single CE packet, the congestion control algorithms followed at the end-systems MUST be essentially the same as the congestion control response to a *single* dropped packet” with “MUST” → “SHOULD” due to ongoing experiments with ECN/AQM
 - RFC 6679: “sender SHOULD mark packets as ECT(0) unless the receiver expresses a preference for ECT(1) or for a random ECT value”
 - Mirja’s discuss: “We are currently discussion mechanisms where the AQM in the congested network node sends much more CE markings than one would see loss (when using TCP) for the same level of congestion. When treating ECN-CE similar to loss, such a different behavior could trigger the circuit breaker unnecessarily”

RTP Circuit Breaker and ECN

- Open questions:
 - Should a circuit breaker respond to ECN-CE?
 - How should a circuit breaker respond to ECN-CE?
 - In general, how should response to ECN-CE relate to response to loss? Is this recommendation changing and, if so, what impact does this have on backwards compatibility?

Should a circuit breaker respond to ECN?

- TSVWG circuit breaker allows ECN to be included in measure of congestion level:
 - “An egress meter can also count ECN [RFC3168] congestion marks as a part of measurement of congestion, but in this case, loss **MUST** also be measured to provide a complete view of the level of congestion.”
- Says nothing on whether loss *has to* occur before circuit breaker triggers
 - Ought to be acceptable, with appropriate threshold, provided the circuit breaker also responds aggressively to loss
 - If this is not acceptable, why not?

What should the response be?

- For media flows, want circuit breaker trigger before loss occurs
 - Loss and latency due to queue overflow disruptive to user experience – we don't want to push queue to overload
 - Would seem okay to trigger with ECN marking rate, m , proportional to $1/\sqrt{m}$ for classic ECN, or $1/m$ for L4S, with appropriate scaling factor; aggressive circuit breaker if loss does occur
 - What should the scaling factor be?
 - Note: must define scaling factor for congestion control – not ECN specific
- Is this acceptable in principle? If not, why not?

General response to ECN-CE

- In general, how should response to ECN-CE relate to response to loss?
 - If this recommendation is changing, what about backwards compatibility?
 - Is the recommendation changing for ECT(0) as well as ECT(1)?
- What requirements do new AQM schemes place on ECN-CE response?
 - Are they compatible with slowly responsive congestion control algorithms, that aggregate ECN-CE feedback and send in a batch, several RTTs after receipt?